


UNIVERSITY OF
NEWCASTLE



University of Newcastle upon Tyne

COMPUTING SCIENCE

Learning by Doing: Practical Courses in Lightweight Formal Methods
using VDM++

P. G. Larsen, J. S. Fitzgerald and S. Riddle.

TECHNICAL REPORT SERIES

No. CS-TR-992

December, 2006

Learning by Doing: Practical Courses in Lightweight Formal Methods using VDM++

P. G. Larsen, J. S. Fitzgerald, S. Riddle.

Abstract

We describe the design and delivery of two courses that aim to develop skills of use to students in their subsequent professional practice, whether or not they are directly applying formal methods. Both curricula take a "lightweight" approach, emphasising skills in model construction and analysis by test rather than formal verification. The accessibility of the formalism is enhanced by the use of well-established notations (VDM-SL and VDM++); its credibility is established by using examples drawn from industrial projects, and by using an industrial-strength tool set. Examples from the courses are presented and initial results based on student evaluation and examination performance are discussed.

Bibliographical details

LARSEN, P. G., FITZGERALD, J. S., RIDDLE, S.

Learning by Doing: Practical Courses in Lightweight Formal Methods using VDM++
[By] P. G. Larsen, J. S. Fitzgerald and S. Riddle.

Newcastle upon Tyne: University of Newcastle upon Tyne: Computing Science, 2006.

(University of Newcastle upon Tyne, Computing Science, Technical Report Series, No. CS-TR-992)

Added entries

UNIVERSITY OF NEWCASTLE UPON TYNE
Computing Science. Technical Report Series. CS-TR-992

Abstract

We describe the design and delivery of two courses that aim to develop skills of use to students in their subsequent professional practice, whether or not they are directly applying formal methods. Both curricula take a "lightweight" approach, emphasising skills in model construction and analysis by test rather than formal verification. The accessibility of the formalism is enhanced by the use of well-established notations (VDM-SL and VDM++); its credibility is established by using examples drawn from industrial projects, and by using an industrial-strength tool set. Examples from the courses are presented and initial results based on student evaluation and examination performance are discussed.

About the author

Peter Gorm Larsen is Professor of Computer Technology and Embedded Systems at The Engineering College of Aarhus, Denmark. He also works as an independent consultant and holds a Visiting Lectureship at Newcastle.

John Fitzgerald is Reader in Computing Science (Dependability) at Newcastle University. He is a specialist in the engineering of dependable computing systems, particularly in rigorous analysis and design tools. His research interests are in system modelling and formal reasoning, reconfigurable and fault tolerant systems.

Steve Riddle is a Lecturer in Computing Science at Newcastle. His research interests are in models of complex systems, traceability and refinement. He leads work in this area for the BAESYSTEMS Dependable Computing Systems Centre at Newcastle.

Suggested keywords

FORMAL METHODS,
SOFTWARE ENGINEERING EDUCATION

Learning by Doing: Practical Courses in Lightweight Formal Methods using VDM⁺⁺

Peter Gorm Larsen¹, John S. Fitzgerald², and Steve Riddle²

¹ Engineering College of Aarhus, Dalgas Avenue 2, DK-8000 Aarhus C, DK

<http://www.iha.dk>, Email: pgl@iha.dk

² School of Computing Science, Newcastle University, NE1 7RU, UK

<http://www.cs.ncl.ac.uk>, Email: Firstname.Lastname@ncl.ac.uk

Abstract. We describe the design and delivery of two courses that aim to develop skills of use to students in their subsequent professional practice, whether or not they are directly applying formal methods. Both curricula take a “lightweight” approach, emphasising skills in model construction and analysis by test rather than formal verification. The accessibility of the formalism is enhanced by the use of well-established notations (VDM-SL and VDM⁺⁺); its credibility is established by using examples drawn from industrial projects, and by using an industrial-strength tool set. Examples from the courses are presented and initial results based on student evaluation and examination performance are discussed.

1 Introduction

Why do we teach formal methods? The majority of graduates from university computer science courses may not use formal methods directly in their subsequent professional practice. However, if they are to make the most of advances in software technologies, and if they are to have the skill to evaluate new technologies and tools as they emerge, they require an appreciation of the principles of abstraction and rigorous reasoning that underpin them. It follows that the vast majority of students in computer science and software engineering should be exposed to formal methods in some form. However, there are some impediments to this. First, there is a perception that formal methods are a heavyweight technology, requiring advanced mathematical skills and demanding an investment of learning effort far in excess of the potential returns. Second, there is a view in the profession as well as among students that these techniques have not been applied seriously, that they are a branch of theory and that they lack credibility.

Our experience of applying formal methods in industry (e.g. [1]) led to the view that lightweight application of formal methods could be cost-effective. Lightweight approaches have been advocated for some time [2,3]. They emphasise carefully targeted formal modelling as an adjunct to other development processes, rather than a wholesale replacement for them. They also emphasise the tractability of formal specification languages over their expressiveness and aim for models that cover significant parts of systems rather than aiming for comprehensiveness. They also open the possibility that partial analysis might be preferred to fully formal verification. Lightweight methods and tools in this sense provide at least some of the benefits of formalism without requiring the full application of highly specialised technology.

This paper describes experiences developing and delivering courses that endeavour to equip students with generic skills of abstraction and rigorous analysis by means of lightweight formal methods using VDM¹ and its support tools. We concentrate on two introductory 5 ECTS credit² courses offered at undergraduate or Masters level in different institutions: the Engineering College of Aarhus (IHA), Denmark, and Newcastle University, UK. Both courses are designed to be accessible to the majority of Computer Science or IT students and both emphasise the credibility of formal methods by using industrial examples and extensive practical work. This emphasis on practice carries through to the assessment techniques used in the courses.

We describe the institutional context and the capabilities of students at whom our courses are targeted (Section 2). Responses to the challenges of accessibility and credibility embodied in our courses are described in Sections 3 and 4. We describe the structure and delivery of the courses (Section 5) and evaluate their current status (Section 6), considering the improvements and modifications we would like to make in future. Although the paper concentrates on introductory courses, we briefly describe the more advanced courses that follow on from these. We conclude with a discussion of the role of lightweight formal methods in the computing curriculum.

2 Context

In order to explain the approach taken to skills development in our courses, it is necessary to understand a little of the institutional contexts as well as the background of the students taking our courses and the experience of the faculty.

2.1 Institutional Context

IHA is a small college specialising in engineering disciplines, including civil, construction, mechanical and electronic engineering, as well as computer technology and embedded systems. The college has particularly strong links with Danish industry, and consequently the educational focus is on applied research aimed at improving industrial competitiveness. The Bachelor of Science (BSc) studies in the IT area have a focus on embedded systems. In collaboration with Aarhus University a Masters programme is provided and it is here that the VDM courses are given.

Newcastle University is one of the UK's smaller research-intensive universities, and one of the first to offer undergraduate degrees in Computing, producing its first graduates in 1969. The majority of students in the School of Computing Science are now undergraduates studying for the BSc honours degree in Computing Science. The School's strong research ethos means that the curriculum is influenced both by the research interests of the faculty as well as by the needs of the industry and the student market. Historically, the school has had a strong practical rather than theoretical focus in its programmes, stressing software design, programming and fault tolerance rather than formal methods of defect avoidance.

¹ The Vienna Development Method – see <http://www.vdmportal.org>.

² ECTS is the European Credit Transfer and Accumulation System. A typical full-time student will study 60 ECTS credits in one academic year.

In different ways, both institutional settings called for formal methods courses that were primarily practical. In IHA, the imperative was to equip students with generic skills for industry application. At Newcastle, the need was for formal methods courses that were consistent with a degree programme and research ethos that was stronger on systems development than on pure computing theory.

2.2 Student Backgrounds

The IHA course is taken as an option by MSc students just after BSc level. Students may go on to take a further specialised course in modelling and analysis of embedded systems. A typical class size is around 15.

In general the background for the IHA students is very practical. They have studied and used C++ and traditional object-oriented paradigms in their BSc studies. However, they have not been exposed to functional programming languages. Their skills in discrete mathematics are quite limited, although some of them have taken another optional 5 ECTS credit course introducing propositional and predicate calculus, relations and functions, recursive structures and induction.

The Newcastle course is taken by all students in the fourth semester of a 6-semester BSc honours programme in Computing Science. The class size is around 60.

When they encounter their first formal methods course, Newcastle students have at least what Boute describes as “Basic Computing Engineering Mathematics” [4]. We rely on basic skills in formal propositional and predicate calculus, relations and functions, recursive structures and induction, in addition to some continuous mathematics. Students have studied and used Java together with the background on the object-oriented paradigm that Java requires, as well as basic program design and algorithm analysis. They have also been exposed to functional programming (in our case, using Haskell); in common with Pepper [5], we found this was very beneficial for students’ practical work in formal modelling.

It is fair to say in both institutions that, when they encounter our courses, students have had much more experience at programming than they have had at conscious abstraction or rigorous reasoning. Indeed, the concept of abstraction is still quite novel, and hence understanding abstraction is a major “knowledge outcome” for our courses. Other outcomes include: understanding the value of rigorous system description; knowledge of relevant analysis and validation methods such as animation, testing, counterexamples and, to a limited extent, formal reasoning.

2.3 Prior Experience with VDM

VDM [6,7] is a well established formal method which has been taught in specialist courses for over twenty years. In its current form, it consists of a formal modelling language VDM-SL, standardised in 1996 [8], a proof theory [9] and a refinement theory [6]. A comparative study of industrial developments with and without formal techniques [1] encouraged us to believe that a lightweight model-oriented formalism could be embedded successfully in an industrial development process provided:

- the formalism is treated as a precise modelling language, a tool for use at levels of abstraction determined by the user, not as part of a greater (refinement-based) methodology;
- the formalism is presented in an accessible way; tools should link to existing tool support and not require users to be trained on platforms specifically to support the formalism;
- tool support is powerful but lightweight; always favouring automation over completeness.

Our experience using this approach, developing relevant industrial training courses and the first courses prepared for Newcastle in the mid-1990s, led to the approach advocated in a text in 1998 [7]. Subsequent experience with industrial applications led to the extension of VDM to support object-oriented design and concurrency, producing VDM⁺⁺ [10]. Further subsequent developments aim to provide a capability for modelling real-time distributed systems [11].

To summarise, the challenge at IHA and Newcastle is to develop courses that are suitable for students with limited mathematical experience but good programming background, within practice-oriented degree programmes. In particular, we must address the issues of accessibility and credibility of formal methods. In the following sections, we describe how the courses offered so far attempt to address these challenges, using examples drawn from the course materials.

3 Enhancing the Accessibility of Formal Methods

3.1 Choice of Formalism

VDM-SL and VDM⁺⁺ are both *model-oriented*. Data is expressed in terms of base types such as numeric types, structureless tokens and enumerations. Structured types may be built from these basic ones using type constructors which include collections such as sets, sequences and mappings. Distinguished instance variables model persistent data if required. Invariants can be expressed over types and over instance variables. Functionality is defined in terms of referentially transparent functions over the defined data types, or operations over the instance variables. Abstraction is provided in data by the unconstrained character of the basic types and type constructors: for example, there is no maximum integer, and sets have unconstrained, though finite, cardinality. Functions and operations may be specified implicitly in terms of preconditions and postconditions, or explicitly in terms of expressions and statements.

Several factors in the choice of the formalism are relevant to its use in teaching. First, the structure of models is familiar to students already versed in object-oriented design. Second, the use of referentially transparent functions is familiar from functional programming. It is difficult to engage students with a language that is not perceived as mainstream [12]. Using a formalism with some familiar features allowed us to focus attention on the less familiar aspects: the abstraction afforded by types and type constructors, specification by means of pre and post-conditions and key issues such as the use of partial operators. Following this, we are able to introduce some elements of rigorous analysis, in particular the generation of proof obligations. These are perhaps best illustrated by means of an example taken from the course material.

Example: The Trusted Gateway This example is derived from a model developed as part of the ConForm project with British Aerospace (Systems and Equipment Ltd) [1]. We will briefly describe the original model, before discussing how it was simplified in order to provide teaching material.

The gateway implements a security policy on communications, in order to prevent secret messages from entering computing systems which are not authorised to process them securely. Messages arrive at the gateway's input port, to be assessed to determine their security classification. High-security messages are passed to a high-security output port, and low-security messages to a low-security output port. message classification is based on whether the message contains special *marker strings* from sets known as *categories*. A state-based model was originally considered which would model the classification of messages one at a time, and a useful abstraction was found in which the messages were arranged into an input sequence.

In our approach, a key principle of abstraction [7] is the removal of detail that is not relevant to a model's purpose. When considering how to use this example in teaching, the question of the model's purpose changes subtly. While remaining a credible model of a trusted gateway, it must additionally provide a non-trivial, representative example of the use of relevant abstractions. At the point in the course where this model is used, the students are taught about sequences, their operators and how to write functions to manipulate them. The model should therefore illustrate the main issues with the use of sequences, which include:

- appropriate choice of sequence datatype;
- sequence definition through enumeration and subrange;
- the relation between sequences and sets, through `elems` and `inds` operators;
- sequence operators head, tail, index and dangers inherent in partial operators;
- how and when to use datatype invariants;
- explicit functions and definition of preconditions.

An extract from the original model is shown in Figure 1. This uses sequence operators such as indexing and concatenation and partial operators (`hd,tl`) in the `Gateway`. However even this representation can be simplified without reducing its utility for teaching. Since the set abstraction has already been dealt with through separate examples, the use of a set of strings to represent a category is not strictly necessary for the purpose of the model: instead, we simplify “high-security” to mean that a string contains the word “SECRET” and “low-security” to mean that a string contains the word “UNCLASSIFIED” (without the word “SECRET”). This allows us to simplify the `Classify` function to merely look for occurrences of the single trigger words, which still achieves the aim of illustrating sequence indexing. A classroom exercise can then be added, extending the model to deal with the category set.

In presenting the example the basic datatypes are described first, with discussion to motivate the choice of representations such as quote types for the classifications and the invariants on `Message` and `String`. The functions to recursively traverse the message stream and classify the input are then introduced in a piecewise fashion. The example is revisited in a later part of the course when validation is considered, by examining the domain-checking proof obligation applied to the `Gateway` function.


```

types
String = seq of char
inv s == s <> [];

Classification = <HI> | <LO>;

Category = set of String;

Message = String
inv m == len m <= 100;

Ports :: high: seq of Message
       low: set of Message;

functions
Gateway: seq of Message * Category -> Ports
Gateway(ms, cat) ==
  if ms = []
  then mk_Ports([],[])
  else let rest_p = Gateway(tl ms, cat)
       in ProcessMessage(hd ms, cat, rest_p);

Classify: Message * Category -> Classification
Classify(m, cat) ==
  if exists hi in set cat & Occurs(hi,m)
  then <HI> else <LO>;

ProcessMessage: Message * Category * Ports -> Ports
ProcessMessage(m, cat, ps) ==
  if Classify(m,cat) = <HI>
  then mk_Ports([m]^ps.high,ps.low)
  else mk_Ports(ps.high, [m]^ps.low)

```

Fig. 1. Trusted Gateway Model

3.2 Role of Tool Support

VDM's recent successes in industrial application have been closely tied to its tool support, and the courses reflect this. Alongside the development of the denotational semantics given in the ISO VDM Standard, a toolset that implemented an operational semantics was developed [13,14]. The product that ultimately resulted from this work, VDM-TOOLS, aimed for cost-effective industrial utility rather than expressive completeness. The VDMTOOLS product is now maintained on a commercial basis by CSK Systems.

VDMTOOLS support syntax and type checking for the full VDM⁺⁺ language, and contains an interpreter for test-based analysis of specifications written within an executable subset. Testing is further supported by coverage analysis tools. The interpreter has a dynamic link library feature allowing external (non-VDM) code to be incorporated. Automatic code generation to C++ and Java are also supported. VDMTOOLS do not (yet) contain direct proof support, although automatic proof obligation generation is available and automated discharging of proof obligations (via HOL) has been demonstrated [15]. A significant feature of VDMTOOLS is the inclusion of a bidirectional link to the Rose tool for UML modelling. In particular, UML class diagrams can be linked directly to object-oriented VDM⁺⁺ models so that changes to the model can be reflected directly in changes to the diagram and vice versa.

The main tool features used in our introductory courses are syntax and type checking, the interpreter and test coverage tools. For example, for the trusted gateway, students use VDMTOOLS to interpret the model with actual test values. The coverage of the tests selected can then be analysed automatically. In addition, students can automatically generate proof obligations and determine whether for each case they have, for example, properly guarded the use of partial operators with pre-conditions.

The use of the interpreter naturally favours executable specifications. Not all models should be written to be directly executable [16], and even underspecified functions may be executed if putative result values are provided. However, for our students, we agree with Utting and Reeves [17] that the direct execution of a model encourages its exploration. Speed is often a problem with executable formal models, but we are fortunate that the VDMTOOLS interpreter has been greatly enhanced in recent years to support its deployment on large industry-scale test sets.

4 Enhancing the Credibility of Formal Methods

In order to better engage our students with formalism, we felt that it was essential to motivate study with realistic examples. This serves to give students confidence in the industrial applicability of formal techniques. Simple examples such as stacks and dictionaries have an important role in introducing particular technical issues, but they do not provide the students an appreciation of the ability to apply the technique for real system development. Every major abstraction concept in our courses is immediately illustrated on an example derived from a real industry case study. Although we often have to simplify the case study so that it is not cluttered with irrelevant material, we can at least still provide information about the original industrial application.

Example: A Congestion Warning System To illustrate our approach to the use of industrial-scale examples, we consider a traffic congestion warning system (CWS). This model originates from ongoing work carried out in The Netherlands by a specialist information systems consultancy in collaboration with the Dutch government transport ministry. The full example is described in VDM⁺⁺ in [10]; here we concentrate on its use within the course at IHA.

A CWS (Fig. 2) warns drivers of areas of high traffic density, suggesting that they reduce speed to minimise the risk of collision. Data from a variety of sensors, including video, loop detectors, and even human observers, is used to determine whether congestion has occurred. Signals are sent to actuators such as roadside signs to convey information to drivers. Such signals obviously need to be consistent and coherent.

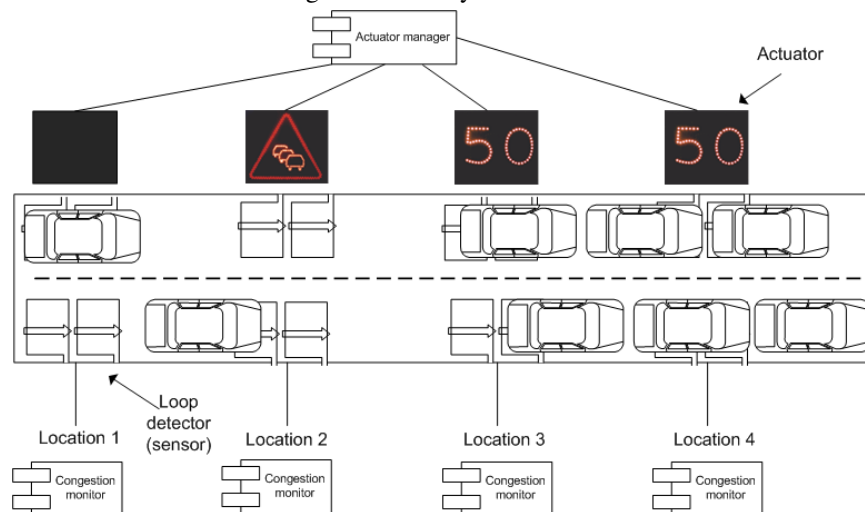


Fig. 2. Overview of the CWS system

In lectures at IHA, a CWS system model is built in a top-down fashion during lectures about sequences and mappings. The overall structure of the CWS is developed in UML (an example class diagram is presented in Fig. 3). In our courses, UML class diagrams are treated as graphical overviews of textual VDM⁺⁺ classes. For example, a fragment of the CWS class from Fig. 3 is shown in Fig. 4. Note how the ordered associations in the UML class diagram are represented as instance variables in the VDM⁺⁺ class and how additional invariants can be constructed, ranging over multiple associations. In practice, the VDMTOOLS-Rose link can be used to maintain consistency between these two views as either is updated [10].

As with the trusted gateway example, the CWS case has been adapted for teaching purposes. In class we recreate some elements of the industrial environment, including changing the requirements and updating the VDM⁺⁺ model accordingly. For example, we introduce a requirement to deal with multiple lanes on a highway, requiring the student to move from modelling using sequences to the use of mappings instead.

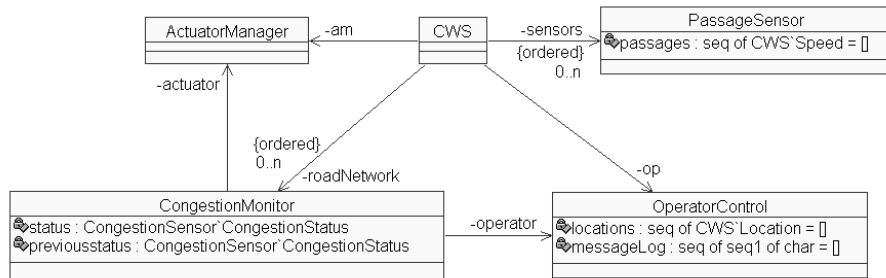


Fig. 3. UML class diagram for CWS

5 Course Aims, Structure and Delivery

5.1 Course Aims and Syllabuses

The introductory courses at IHA and Newcastle aim to develop students' knowledge and skill in building and analysing suitably abstract models of computing systems. In addition, we aim to develop students' understanding of *why* formal models should be built and the issues involved in commercial applications of formal techniques. Basic modelling principles are taught using a combination of UML and VDM⁺⁺.

The IHA and Newcastle courses have similar structures. Following initial motivation, the elements of the modelling language are rapidly introduced in overview. The key abstractions are then introduced, starting with basic data and functional modelling, then moving on to the main collection abstractions (sets, mappings and sequences). Along the way, additional functional abstraction techniques including implicit (pre/post) specification and data type invariants are introduced. Each abstraction is illustrated on an industry-inspired example like the trusted gateway and CWS models described above. These presentations use material from Chapters 1-8 of [10]. Although the focus is on modelling, rigorous approaches to analysis are introduced via the proof obligations.

```

class CWS
...
instance variables
roadNetwork: seq of CongestionMonitor := [];
sensors      : seq of PassageSensor := [];
inv len roadNetwork = len sensors;

end CWS
  
```

Fig. 4. Congestion Warning System: extract

5.2 Delivery Methods

Both courses are presented over a 6-7 week period within one semester. At Newcastle, this amounts to about 24 hours of lectures and 12 hours of practical laboratory classes (with teaching support) in addition to private study. The IHA course³ involves 24 hours of lectures and 18 hours of presentations by each group of students reporting on their practical work.

The difference in class sizes has to be taken into account in considering the delivery methods at each institution. In both IHA and Newcastle, we aim for a high level of interaction between lecturer and students. Typically, practical questions are presented on the slides immediately after a subject has been introduced. In addition, at IHA, with a smaller class size, the students know that in effect they will take turns in answering these questions. As a consequence they need to stay alert and active to be able to answer the next question they get. This has the consequence that the students learn the new concepts over a longer period than if they just prepare for the final examination. We believe that learning the topics in this fashion increases the retention rate.

At both IHA and Newcastle, students carry out practical coursework using VDM-TOOLS, to which they are introduced at a very early stage. Here they need to learn abstraction “by doing” it themselves on their projects. Although the projects carried out by the students are not large, they are all inspired by real systems and thus have more of a flavour of realism. This coursework element is particularly strong in the IHA course, where it forms the main part of the assessment. At IHA, students are offered a list of suggested projects. The current selection includes:

- SAFER (Simplified Aid For Extra Vehicular Activity Rescue) from [18];
- the production cell from [19];
- a cash dispenser from [20];
- CyberRail from [21];

Students may also select projects from other MSc courses in distributed real-time and IT systems or suggest their own topic. The IHA course is mainly targeted at students studying the development of embedded software-based systems, so the topics are all reactive systems. Students in the most recent cohort selected all of the projects listed above, and three groups suggested their own projects based on a car radio navigation system, a self-navigating vehicle and a personal medical unit. Some of the projects already have substantive VDM models whereas for others only ideas are present and/or models are provided in different formalisms. Those students starting from existing models concentrate on analysing and modifying those models. Students selecting projects without existing models enjoy more freedom but face the challenge of choosing an appropriate level of abstraction.

Coursework at Newcastle accounts for 20% of the total course assessment. The students initially have to complete a short series of tutorial exercises working with basic datatypes and set, sequence and mapping abstractions, and interacting with VDM-TOOLS. In contrast to the choice of projects at IHA, at Newcastle the main exercise requires all students to understand and extend a given model. In recent years the model

³ <http://kurser.iha.dk/eit/tivdml/>

has taken the form of an automated supermarket checkout system, flight cargo management and an organisation's security access control protocol. In each case, students are given a basic model in which some abstraction decisions have already been made, and are asked to extend it by writing functions and extending datatypes. Typically the final part of the coursework will require the student to decide how to model a more complex extension: for example, in the access control model, students were asked to model the Least Privileges Problem. This gives students more opportunity to make their own abstraction decisions. With the large class size, having all students work on the same scenario permits an even level of support to all students, though it does reduce the scope for more innovative ideas from stronger students. The open-ended final question is intended to address this point.

5.3 Assessment

We are not concerned with assessing the students' ability to learn by rote the details of the language syntax: rather, we wish to assess their higher-level understanding of abstraction, rigorous system modelling and reasoning. Our assessment mechanisms for the two courses share this common principle, but realise it in different ways.

The assessment of the IHA course is carried out by means of an oral examination. Students are required to prepare a presentation of no more than 10 minutes, to be delivered without multimedia support. They are also told to expect to be interrupted with questions. Each student is examined for 15 minutes and 5 additional minutes were used to decide upon the grade with the external examiner. The nature and difficulty of the questions asked is tailored to each student on the basis of the report submitted prior to the examination, and during the examination itself. Broadly speaking, the better the written and oral performance, the more demanding the questions. The examination attempts to cover as much of the curriculum as possible both in depth and in breadth. The final grade is an aggregate of the assessments of the written report and the oral examination.

The Newcastle class is about four times the size of the IHA class, so a written examination is used, coupled with assessment of laboratory coursework. Currently, 80% of the final mark for the course is derived from the written examination, 20% from the coursework (to pass, students must obtain passing marks separately in both components). At first glance, this is a very conventional arrangement. However, the examination is quite unconventional. First, it is "open-book", meaning that students may consult printed or written materials. Second, it is based on a scenario that students receive some weeks ahead of the examination. An example scenario is in Fig. 5.

Students have every opportunity to read the scenario, anticipate questions and prepare answers which they can validate through VDMTOOLS. Since the examination is open book, these prepared answers can be consulted during the examination itself. Students are presented with previously unseen questions in which they are asked to comment on, explain, extend and revise the model described in the scenario. This scheme allows us to ask some gentle questions that the students might reasonably have anticipated, but also deeper questions requiring extensions to the model, for example, that we could not ask in a conventional closed-book examination. In our example scenario, basic questions could include asking students to provide comments explaining the data type

Imagine that you are a member of a team commissioned to develop memory management functions for an operating system. In order to gain a better understanding of the management policy to be used, you have been asked to produce a formal model. An initial draft of the model in VDM-SL is shown here. Memory management is used to allocate partitions (contiguous regions of computer memory) to jobs scheduled to run in the system. Each job to be executed has an identifier. This is modelled by the `JobId` type defined as follows:

```
JobId = token;
```

The partitions allocated to jobs have a start location in memory. Locations are modelled as natural numbers. Each job has a certain size which represents the amount of memory required to execute the job. The system maintains two mappings recording job size and memory location. There is also an upper limit on the memory locations handled by the system. These concepts are modelled by the following type definitions:

```
Location = nat;
Size = nat;

JobLoc = map JobId to Location;
JobSize = map JobId to Size;

System :: jobs : JobLoc
         alloc : JobSize
         upperLimit : Size
inv sys == dom sys.alloc subset dom sys.jobs;
```

Jobs are allocated to a particular memory location by the function `allocate`, defined as follows:

```
allocate: System * JobId * Location -> System
allocate(mk_System(jobs, alloc, lim), j, l) ==
    mk_System(jobs, alloc union j |-> l, lim)
pre j not in set dom alloc;
```

There are also further constraints on the system, such as:

1. Jobs allocated must fit within the upper memory limit
2. Jobs must not overlap

Further functionality to be added to the model will include:

- De-allocating a job from memory
- Providing a memory address where a given job will fit
- Managing the scheduling of jobs

Note that neither of these lists is intended to be exhaustive.

Fig. 5. Examination Scenario

invariant on `System`, or formalising the constraints. After asking students to explain the purpose of the precondition on the `allocate` function, we can ask if the function respects the invariant (checking understanding of proof obligations) and to modify it appropriately. Many of these questions could be anticipated by a well prepared student. The more challenging question would ask the student to extend the model with specifications of priority-based queueing for jobs, for example.

6 Experiences, Evaluation and Results

The IHA course has been offered just once, while the Newcastle course has been running for three years in its current form, and for about ten years in previous incarnations. Both institutions gather student feedback via anonymous questionnaires, and both ask students to assess the relevance of the course to the degree programme, and the difficulty. Evaluations at both institutions have been encouraging. Student evaluation is high on relevance to the degree programme (all students rating this good or very good at IHA; 80% of students rating this satisfactory or high at Newcastle). There is a broader spread of opinion on difficulty, about a third of IHA students and around half of Newcastle students finding their course challenging. Bearing in mind that the Newcastle course is compulsory, this evaluation is broadly typical of Computing Science courses at present. In their free comments, students have been particularly positive about clear and precise teaching (some finding it too slow), and about the practical tool-supported approach taken. We have had few comments on the relevance or irrelevance of the material.

The practical projects carried out by the students were appropriate for the students to be able to acquire experience producing and analysing abstract VDM⁺⁺ models independently. Projects building on a pre-existing VDM model were less challenging because an appropriate level of abstraction had already been found. This kind of project was most appropriate for the weaker students. On the other hand the projects where no prior model existed gave students the potential to explore and develop the project to ensure greater coverage of the range of modelling elements and techniques covered in the course.

At IHA the students in the first cohort performed remarkably well at the examination, well above average compared with other courses. In Newcastle the performance by the students was typical for compulsory courses.

Bearing in mind the student evaluations and performance in assessments, there are several changes that we will consider for the next delivery of both courses. At IHA, no specific training was given in the use of VDMTOOLS and Rational Rose. It turned out that a little more practical introduction to VDMTOOLS, as provided at Newcastle, would have been an advantage. Our plans for the next academic year include conducting a few small practical exercises with VDMTOOLS together with the students during the first week of the course to accommodate this.

During the IHA course, strict progress checks on the group projects were imposed, with one presentation on progress made by each group per week. Communication between the different groups was facilitated by requiring the groups to give each other feedback. Unfortunately some of the reports are in Danish so it will only be partly use-

ful for international teachers. We will consider in the future to make it mandatory to all students writing their reports in English.

Higher-level courses IHA students have a further optional course on embedded and real-time systems building on their introduction to formal modelling. This runs for the first time in October 2006 and includes:

- Model Structuring and Combining Views (Chapters 9 and 10 from [10])
- Concurrency in VDM⁺⁺ (Chapter 12 from [10])
- Real-time modelling in VDM⁺⁺ over two weeks
- Distributed systems in VDM⁺⁺
- Model Quality (Chapter 13 from [10])

The real-time and distribution subjects will be taught using new primitives added to VDM⁺⁺ in current research [22,11]. It will be interesting to see how this develops, since the students will most likely be the first real users of the new tool support.

At Newcastle, after completing the introductory course, some students go on to study a further optional 10 ECTS credit course in their final year which focuses on the achievement of *correctness*, via a spectrum of formal and informal, static and dynamic approaches. This course builds on students' abstraction and modelling skills to gain advanced skills in formal verification, refinement, annotation-based design by contract, and static analysis. This latter course is not focussed on a single formalism, but uses Hoare Logic, VDM for refinement and Java Markup Language for static analysis.

7 Concluding Remarks

We have presented a practical approach to the lightweight use of an established formal method as a means of developing abstraction and rigorous analysis skills among Computing Science students. The principles underpinning the approach relate to the use of accessible formal notation, supported by strong tools and introduced via examples derived from industrial practice. The approach has been realised in two introductory courses in separate institutions, within slightly different institutional frameworks in small and large classes, in optional and compulsory courses. We believe that a pragmatic approach to introducing lightweight use of formal techniques with plenty of hands-on experience is a way to stimulate interest in this technology and to establish its credibility.

Our university courses have their origins in industry training. Such training is improved if the methods taught are rapidly applicable [23]. In the academic context, the focus must be more on deeper long-term skills that are independent of the particular formalism, but we believe that this has to be balanced with the need to encourage student engagement by using an established, tool-supported formalism.

Student motivation is important. As Reed and Sinclair [24] have suggested, many students are driven by the need to develop skills useful to them in a subsequent career, rather than a fascination for computing science's beauty. Many, in our experience, are also driven by the satisfaction that comes from creating good software and seeing it

run. Students may see formal methods as an obstacle to this. Our lightweight use of formalism is an attempt to deal with this tension.

We agree with Boute [4] that the concept of lightweight formal methods may be mis-used as a means of lowering the threshold of mathematical skill needed to develop good computing systems. We make a point of requiring specific mathematical skills. Indeed, we would argue that the formalism used is not compromised in our courses⁴. By “lightweight” we mean sharply focussing the formalism on particular system concerns (abstraction), in the manner suggested by Jackson and Wing [3].

We may revisit our opening question: why do we teach formal methods? Our response was that the study of formal methods develops skills, notably those of abstraction and rigorous analysis, that will be of use beyond the immediate course in which they are taught. But are our students better at abstraction in general as a result of taking our courses? We believe that few teachers could answer this question convincingly. In future work, we aim to devise an evaluation, independent of the particular formalism, that will help us determine the extent to which formal methods skills are truly “transferable”.

Acknowledgments We are grateful to Erik Ernst, Jozef Hooman, Troels Fedder Jensen, Marcel Verhoef, Stefan Wagner, Jeremy Bryans and referees for their valuable comments.

References

1. Larsen, P.G., Fitzgerald, J.S., Brookes, T.: Applying Formal Specification in Industry. *IEEE Software* **13** (1996) 48–56
2. Jones, C.B.: A Rigorous Approach to Formal Methods. *IEEE Computer* **29** (1996) 20–21
3. Jackson, D., Wing, J.: Lightweight Formal Methods. *IEEE Computer* **29** (1996) 22–23
4. Boute, R.T.: Can Lightweight Formal Methods Carry the Weight? In Duce, D.A., et al., eds.: *Teaching Formal Methods: Practice and Experience 2003*. Oxford Brookes University (2003) Available at <http://cms.brookes.ac.uk/tfm2003/>.
5. Pepper, P.: Distributed Teaching of Formal Methods. In Dean, C.N., Boute, R.T., eds.: *Teaching Formal Methods*. Volume 3294 of *Lecture Notes in Computer Science*. Springer (2004) 140–152
6. Jones, C.B.: *Systematic Software Development Using VDM*. Second edn. Prentice-Hall International, Englewood Cliffs, New Jersey (1990)
7. Fitzgerald, J., Larsen, P.G.: *Modelling Systems – Practical Tools and Techniques in Software Development*. Cambridge University Press, The Edinburgh Building, Cambridge CB2 2RU, UK (1998) ISBN 0-521-62348-0.
8. Andrews, D.J., ed.: *Information technology – Programming languages, their environments and system software interfaces – Vienna Development Method – Specification Language – Part 1: Base language*. International Organization for Standardization (1996) International Standard ISO/IEC 13817-1.
9. Bicarregui, J.C., Fitzgerald, J.S., Lindsay, P.A., Moore, R., Ritchie, B.: *Proof in VDM: A Practitioner’s Guide*. FACIT. Springer-Verlag (1994)

⁴ Bernhard Steffen has suggested that we have not introduced a lightweight formal method. The specific weight of VDM remains the same – we have merely provided tools for lifting it!

10. Fitzgerald, J., Larsen, P.G., Mukherjee, P., Plat, N., Verhoef, M.: Validated Designs for Object-oriented Systems. Springer, New York (2005)
11. Verhoef, M., Larsen, P.G., Hooman, J.: Modeling and validating distributed embedded real-time systems with VDM++. In Misra, J., Nipkow, T., Sekerinski, E., eds.: FM 2006: Formal Methods. Volume 4085 of Lecture Notes in Computer Science., Springer (2006) 147–162
12. Paige, R.F., Ostroff, J.S.: Specification-Driven Design with Eiffel and Agents for teaching Lightweight Formal Methods. In Dean, C.N., Boute, R.T., eds.: Teaching Formal Methods. Volume 3294 of Lecture Notes in Computer Science. Springer (2004) 107–123
13. Larsen, P.G., Lassen, P.B.: An Executable Subset of Meta-IV with Loose Specification. In: VDM '91: Formal Software Development Methods, VDM Europe, Springer-Verlag (1991)
14. Elmstrøm, R., Larsen, P.G., Lassen, P.B.: The IFAD VDM-SL Toolbox: A Practical Approach to Formal Specifications. ACM Sigplan Notices **29** (1994) 77–80
15. Agerholm, S., Sunesen, K.: Reasoning about vdm-sl proof obligations in hol. Technical report, IFAD (1999)
16. I.J. Hayes, C.J.: Specifications are not (necessarily) executable. Software Engineering Journal (1989) 330–338
17. Utting, M., Reeves, S.: Teaching Formal Methods Lite via Testing. Journal of Software Testing, Verification and Reliability **11** (2001) 181–195
18. NASA: Formal Methods, Specification and Verification Guidebook for Verification of Software and Computer Systems. Vol 2: A Practitioner's Companion. Technical Report NASA-GB-001-97, Washington, DC 20546, USA (1997) Available from http://eis.jpl.nasa.gov/quality/Formal_Methods/.
19. Lewerentz, C., Lindner, T., eds.: Formal development of reactive systems: case study production cell. Volume 891 of LNCS. Springer-Verlag, New York (1995)
20. Group, T.V.T.: A “Cash-point” Service Example. Technical report, CSK (2005) http://www.vdmportal.org/twiki/pub/Main/VDMPPEexamples/cashdispenser_a4.pdf.
21. RTRI: The Concept of CyberRail. <http://cyberrail.rtri.or.jp/english/> (2006)
22. Group, T.V.T.: Development Guidelines for Real Time Systems using VDMTools. Technical report, CSK (2006)
23. Loomes, M., Christianson, B., Davey, N.: Formal Systems, Not Methods. In Dean, C.N., Boute, R.T., eds.: Teaching Formal Methods. Volume 3294 of Lecture Notes in Computer Science. Springer (2004) 47–64
24. Reed, J.N., Sinclair, J.E.: Motivating Study of Formal Methods in the Classroom. In Dean, C.N., Boute, R.T., eds.: Teaching Formal Methods. Volume 3294 of Lecture Notes in Computer Science. Springer (2004) 32–46